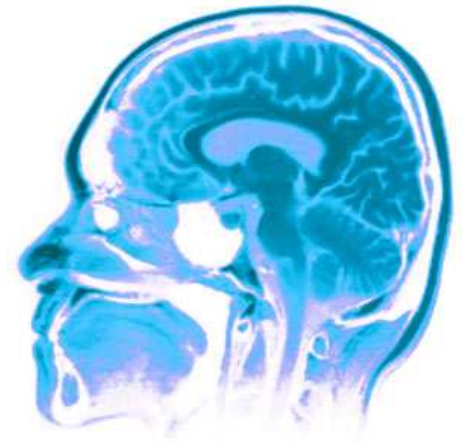




# CPSC540



Linear dimensionality reduction



Nando de Freitas  
*September, 2011*  
*University of British Columbia*

# Outline

We introduce the Singular Value Decomposition (SVD). This is a matrix factorization that has many applications, including:

- Information retrieval,
- Least-squares problems,
- Image processing,
- Dimensionality reduction.

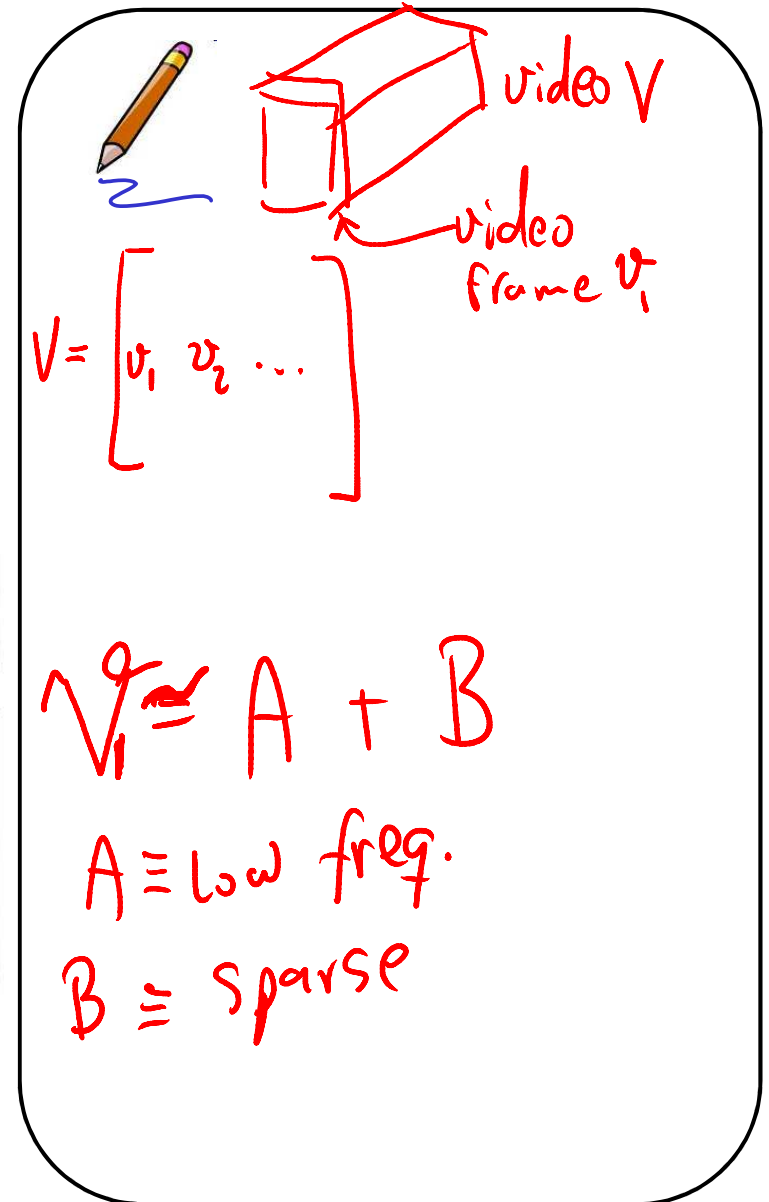
# A video can be treated as a matrix that can be decomposed



(a) Original frames

(b) Low-rank  $\hat{L}$

(c) Sparse  $\hat{S}$



# Decomposing images



(a)  $M$

(b)  $\hat{L}$

(c)  $\hat{S}$

[Candes et al, 2009]

# Eigenvalue decomposition

Let  $\mathbf{A}$  be an  $m \times m$  matrix of reals; that is  $\mathbf{A} \in \mathbb{R}^{m \times m}$ . If we place the eigenvalues of  $\mathbf{A}$  into a diagonal matrix  $\mathbf{\Lambda}$  and gather the eigenvectors into a matrix  $\mathbf{Q}$ , then the eigenvalue decomposition of  $\mathbf{A}$  is given by

$$\mathbf{A} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^{-1}.$$



$$A \mathbf{q}_i = \lambda_i \mathbf{q}_i$$

$$Q = [\mathbf{q}_1 \ \mathbf{q}_2 \ \dots \ \mathbf{q}_m]$$

$$\mathbf{\Lambda} = \begin{bmatrix} \lambda_1 & & & \\ & \ddots & & \\ & & \ddots & \\ 0 & & & \lambda_m \end{bmatrix}$$

If  $A$  is symmetric, then  $A = A^T$

$$A = Q \mathbf{\Lambda} Q^T$$

# SVD decomposition

$$\underline{\mathbf{A} \in \mathbb{R}^{m \times n}}$$

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$$

thin SVD

$\mathbf{\Sigma} \in \mathbb{R}^{n \times n}$  is diagonal with positive entries (singular values in the diagonal).

$\mathbf{U} \in \mathbb{R}^{m \times n}$  has orthonormal columns.

$\mathbf{U} \in \mathbb{R}^{m \times m}$

$\mathbf{V} \in \mathbb{R}^{n \times n}$  has orthonormal columns and rows.

That is,  $\mathbf{V}$  is an orthogonal matrix, so  $\underline{\mathbf{V}^{-1} = \mathbf{V}^T}$ .

$$\mathbf{U} = \begin{bmatrix} \underline{u}_1 & \underline{u}_2 & \dots & \underline{u}_n \end{bmatrix}$$

$$\underline{u}_i^T \underline{u}_j = \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases}$$

# SVD

$$A = U \Sigma V^T$$

The equations relating the right singular values  $\{\mathbf{v}_j\}$  and the left singular vectors  $\{\mathbf{u}_j\}$  are

$$\mathbf{A} \mathbf{v}_j = \sigma_j \mathbf{u}_j \quad j = 1, 2, \dots, n$$

$$\mathbf{A} \mathbf{V} = \mathbf{U} \Sigma$$

$$A = U \Sigma V^T$$

$$AV = U \Sigma V^T V$$

$$AV = U \Sigma$$

$$\mathbf{A} \begin{bmatrix} \mathbf{v}_1 & \mathbf{v}_2 & \dots & \mathbf{v}_n \end{bmatrix} = \begin{bmatrix} \mathbf{u}_1 & \mathbf{u}_2 & \dots & \mathbf{u}_n \end{bmatrix} \begin{bmatrix} \sigma_1 & \sigma_1, \sigma_2, \dots & & \\ & \sigma_2 & & \\ & & \dots & \\ & & & \sigma_n \end{bmatrix}$$

# SVD properties

1. There is no assumption that  $m \geq n$  or that  $\mathbf{A}$  has full rank.
2. All diagonal elements of  $\mathbf{\Sigma}$  are non-negative and in non-increasing order:

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_p \geq 0$$

where  $p = \min(m, n)$



# SVD in terms of eigenvalues

**Theorem 4** *The nonzero singular values of  $\mathbf{A}$  are the (positive) square roots of the nonzero eigenvalues of  $\mathbf{A}^T \mathbf{A}$  or  $\mathbf{A} \mathbf{A}^T$  (these matrices have the same nonzero eigenvalues).*

$$\Sigma = \Sigma^T \\ U^T U = I$$



$$\begin{aligned} \underbrace{A^T A}_{\substack{n \times n \\ n \times n}} &= (U \Sigma V^T)^T (U \Sigma V^T) \\ &= (V \Sigma^T U^T) (U \Sigma V^T) = V \bar{\Sigma} \Sigma V^T \\ &= V \Sigma^2 V^T = Q \Lambda Q^T \end{aligned}$$

$$\begin{aligned} V &= Q \\ \Sigma &= \sqrt{\Lambda} \end{aligned}$$



$$AA^T = U\Sigma V^T V\Sigma U^T = U\Sigma^2 U^T$$

$m \times m$

$$A = U\Sigma V^T$$

$$AV = U\Sigma$$

$$\underline{AV} \underline{\Sigma}^{-1} = U$$

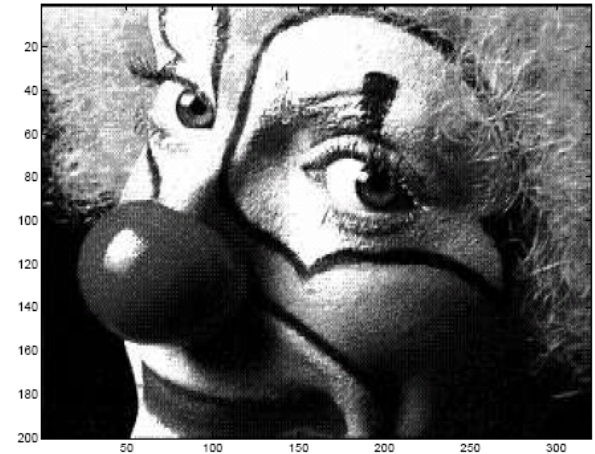
$\rightarrow V =$

$$A = \begin{bmatrix} 1 & 2 & 3 & 4 & 5 \end{bmatrix}$$

# Image compression example in python

```
from scipy import *
from pylab import *

img = imread("clown.png")[:, :, 0]
gray()
figure(1)
imshow(img)
```



```
m,n = img.shape
U,S,Vt = svd(img)
S = resize(S,[m,1])*eye(m,n)

k = 20
figure(2)
imshow(dot(U[:,1:k],dot(S[1:k,1:k],Vt[1:k,:])))
show()
```



# The truncated SVD



$$A = U \Sigma V^T$$

$$m = 5$$

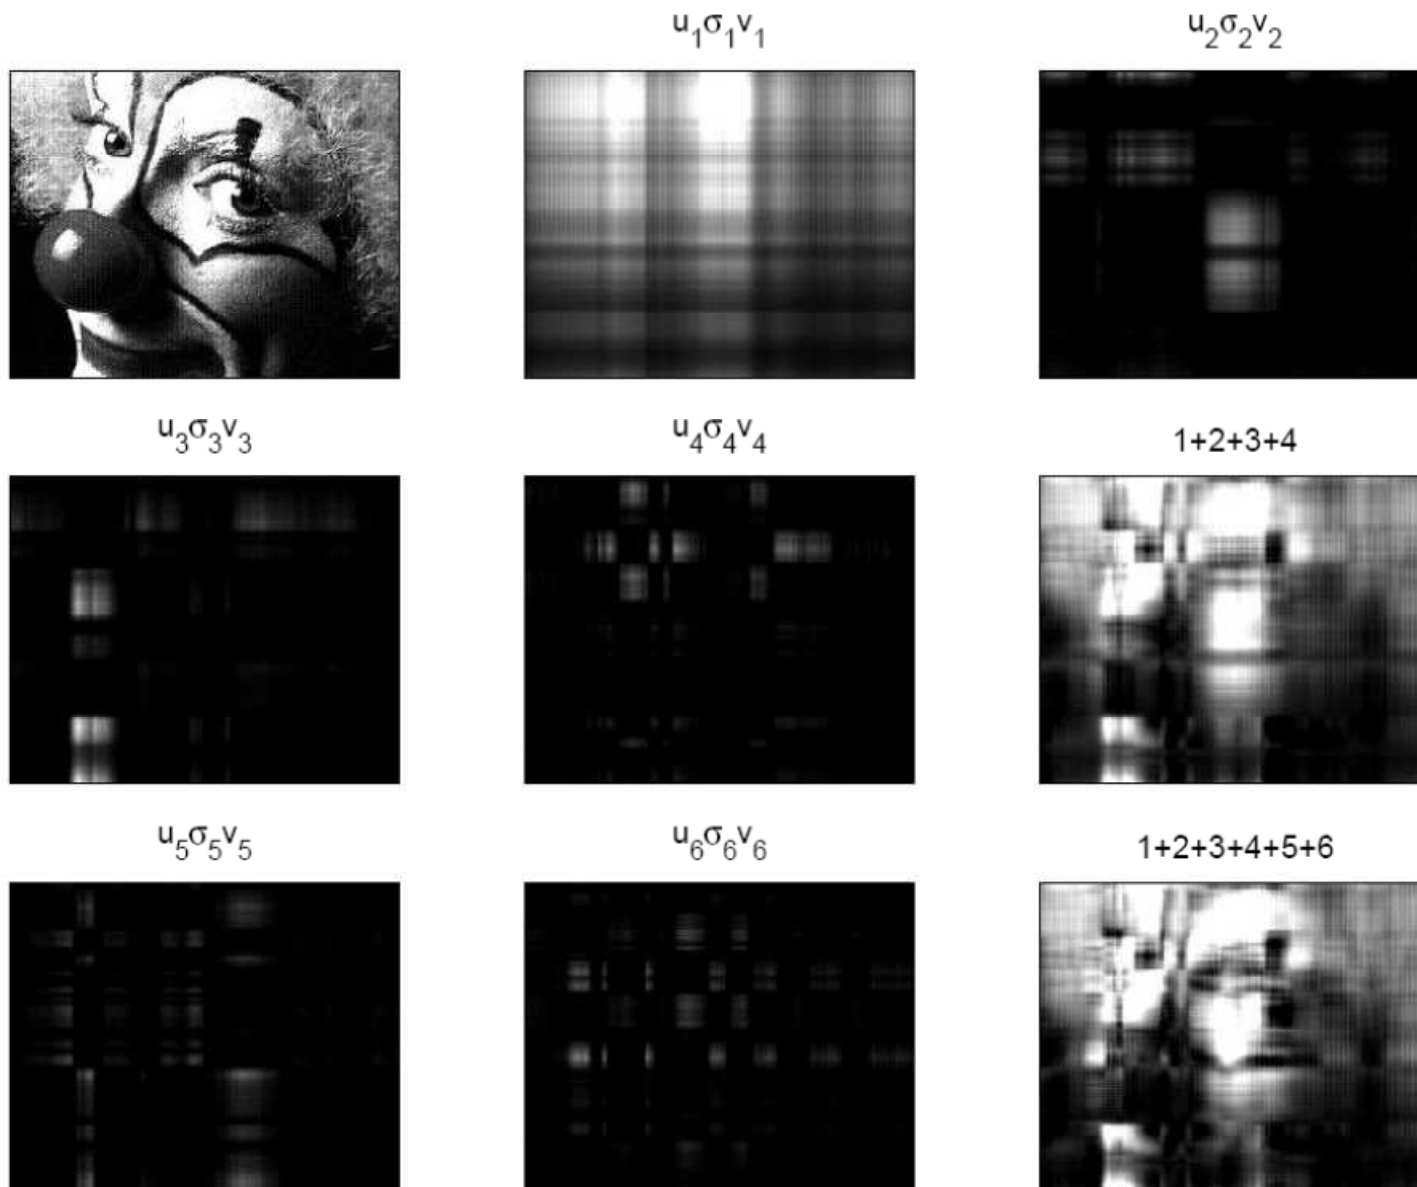
$$n = 4$$

$$k = 2$$

$$A = \underline{u}_1 \underline{\sigma}_1 \underline{v}_1^T + \underline{u}_2 \underline{\sigma}_2 \underline{v}_2^T + \underline{u}_3 \underline{\sigma}_3 \underline{v}_3^T + \underline{u}_4 \underline{\sigma}_4 \underline{v}_4^T$$

$$A = \begin{bmatrix} \underline{u}_1 & \underline{u}_2 & \underline{u}_3 & \underline{u}_4 \\ \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \end{bmatrix} \begin{bmatrix} \underline{\sigma}_1 & 0 & 0 & 0 \\ 0 & \underline{\sigma}_2 & 0 & 0 \\ 0 & 0 & \underline{\sigma}_3 & 0 \\ 0 & 0 & 0 & \underline{\sigma}_4 \end{bmatrix} \begin{bmatrix} \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \end{bmatrix} \begin{matrix} \underline{v}_1^T \\ \underline{v}_2^T \\ \underline{v}_3^T \\ \underline{v}_4^T \end{matrix}$$

$$A_k = \underline{u}_1 \underline{\sigma}_1 \underline{v}_1^T + \underline{u}_2 \underline{\sigma}_2 \underline{v}_2^T ; k=2$$



Smaller eigenvectors capture high frequency variations (small brush-strokes).



$$\text{Image} = u_1 \sigma_1 v_1 + u_2 \sigma_2 v_2 + u_3 \sigma_3 v_3 + u_4 \sigma_4 v_4$$

The diagram illustrates the decomposition of an image into its principal components. On the left is the original grayscale image of a clown's face. This image is equal to the sum of four principal components, each represented by a square image. The components are labeled as  $u_1 \sigma_1 v_1$ ,  $u_2 \sigma_2 v_2$ ,  $u_3 \sigma_3 v_3$ , and  $u_4 \sigma_4 v_4$ . The first component shows the most prominent features, while subsequent components capture progressively finer details and textures. Red plus signs and an equals sign are used to indicate the summation.

# Image compression example

The code:

- loads a clown image into a 200 by 320 array A,
- displays the image in one figure,
- performs a singular value decomposition on A,
- displays the image obtained from a rank-20 SVD approximation of A in another figure.



The original storage requirements for A are:

$$320 \times 200 \quad \text{bytes / pixel} = \text{byte}$$

say

The compressed representation requires:

$$320 \times 20 + 200 \times 20 + 20 < 320 \times 200$$

# Text retrieval:

## Latent semantic indexing (LSI)

The SVD can be used to cluster documents and carry out information retrieval by using concepts as opposed to exact word-matching.

This enables us to surmount the problems of synonymy (car, auto) and polysemy (money bank, river bank).



The data is available in a term-frequency (TF) matrix:

$$A = \begin{matrix} \text{word 1} \\ \text{word 2} \\ \vdots \\ \cdot \end{matrix} \begin{bmatrix} \text{doc 1} & \text{doc 2} & \text{doc 3} \\ 1 & 3 & 2 \\ 0 & 1 & 0 \\ \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \end{bmatrix} \quad a_{ij} = \# \text{ times word } i \text{ appears in document } j$$



# LSI example

$$m = 5$$
$$n = 3$$



$$d_1 = \{I, \text{eat}, \text{chips}\}$$

$$d_2 = \{\text{computer}, \text{chips}, \text{chips}\}$$

$$d_3 = \{\text{intel}, \text{computer}, \text{chips}\}$$

$$A = \begin{array}{ccc|c} & d_1 & d_2 & d_3 & \\ \hline & 1 & 0 & 0 & I \\ & 1 & 0 & 0 & \text{eat} \\ & 1 & 2 & 1 & \text{chips} \\ & 0 & 1 & 1 & \text{Computer} \\ & 0 & 0 & 1 & \text{intel} \end{array}$$

$$A = U \Sigma V^T$$

$$U^T A = U^T U \Sigma V^T$$

$$U^T A = \Sigma V^T$$

$$\Sigma^{-1} U^T A = \Sigma^{-1} \Sigma V^T$$

$$V^T = \Sigma^{-1} U^T A$$

# Truncated SVD for LSI

If we truncate the approximation to the  $k$ -largest singular values, we have

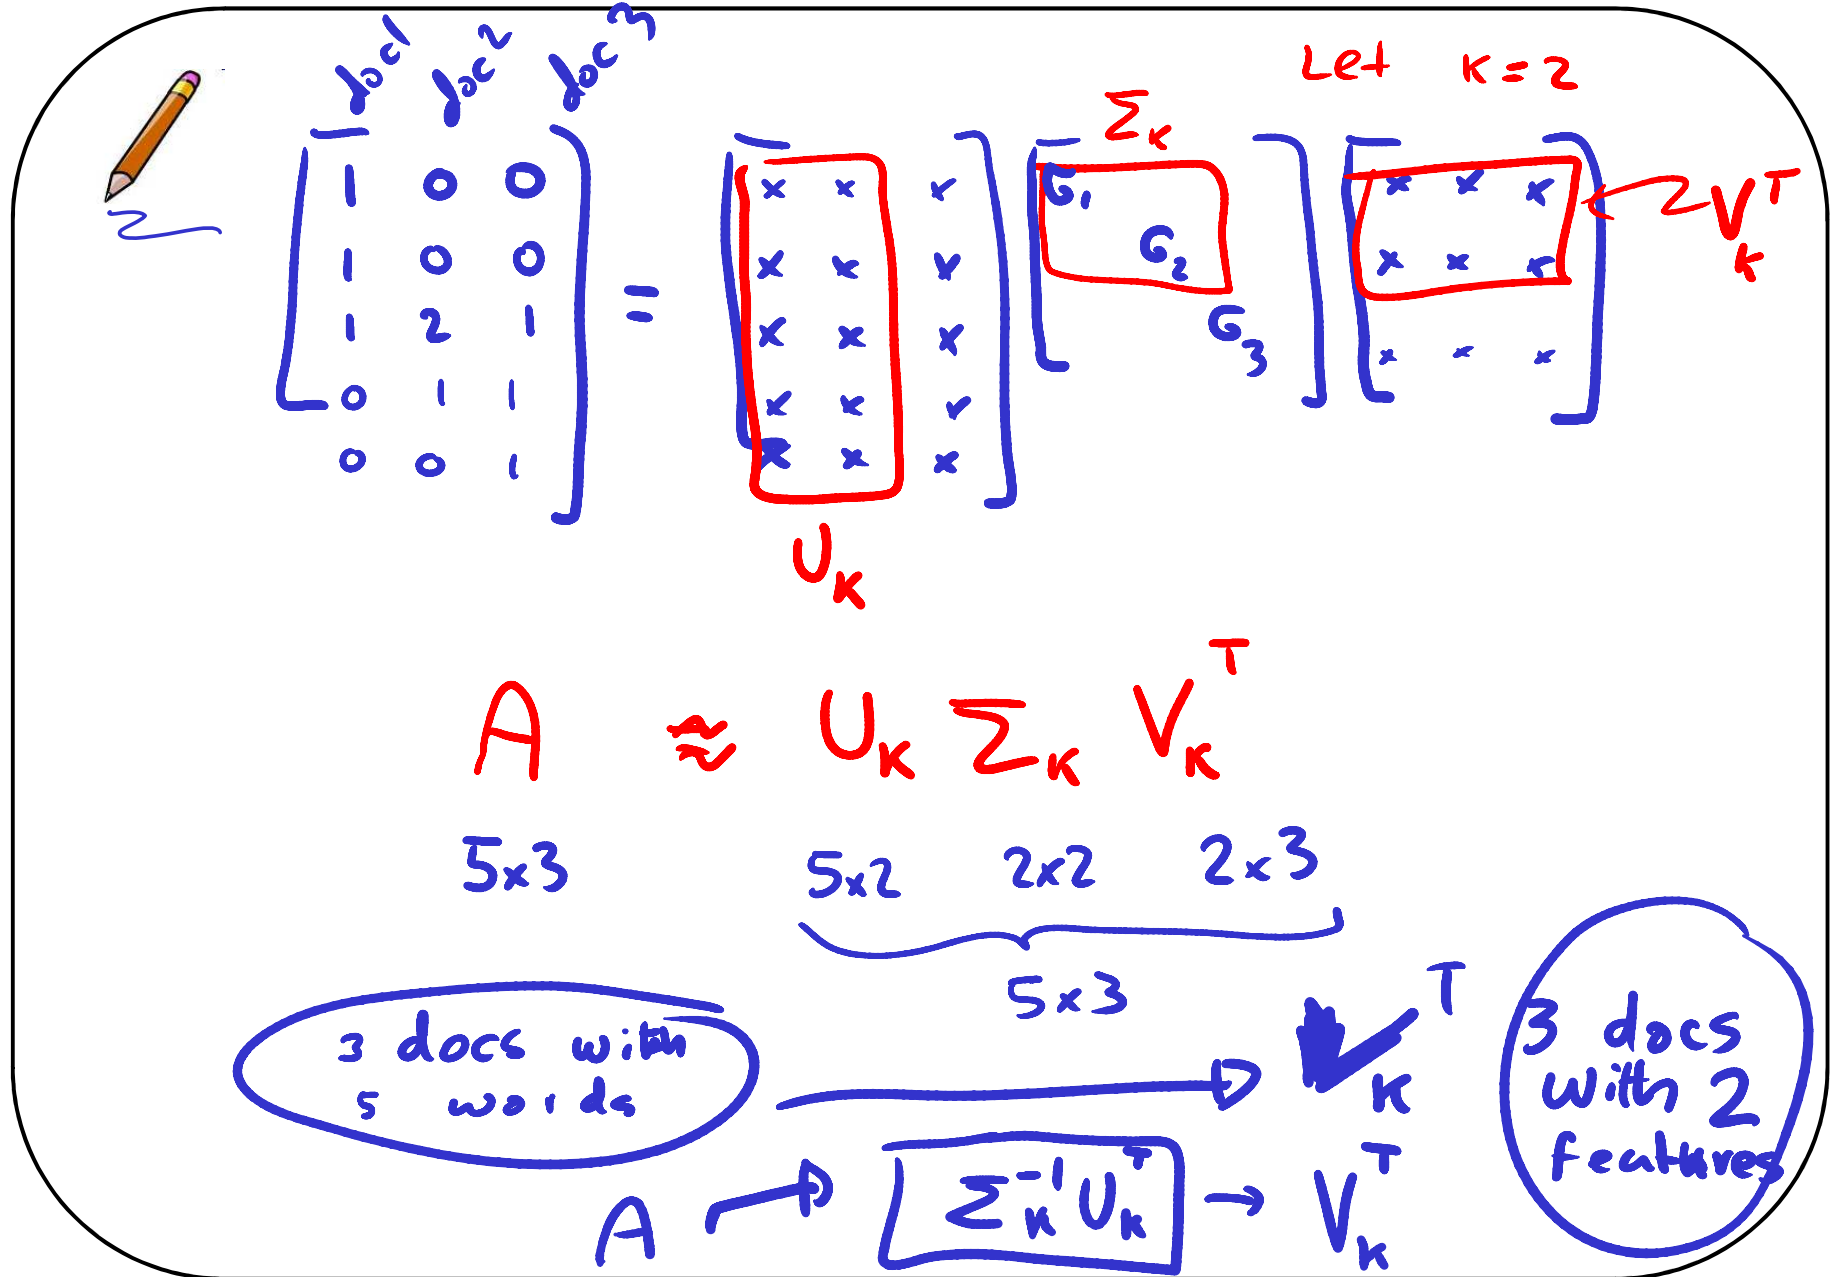
$$\mathbf{A} = \mathbf{U}_k \mathbf{\Sigma}_k \mathbf{V}_k^T$$

So

$$\mathbf{V}_k^T = \mathbf{\Sigma}_k^{-1} \mathbf{U}_k^T \mathbf{A}$$

In English,  $\mathbf{A}$  is projected to a lower-dimensional space spanned by the  $k$  singular vectors  $\mathbf{U}_k$  (eigenvectors of  $\mathbf{A}\mathbf{A}^T$ ).

# Part I: Building the search engine



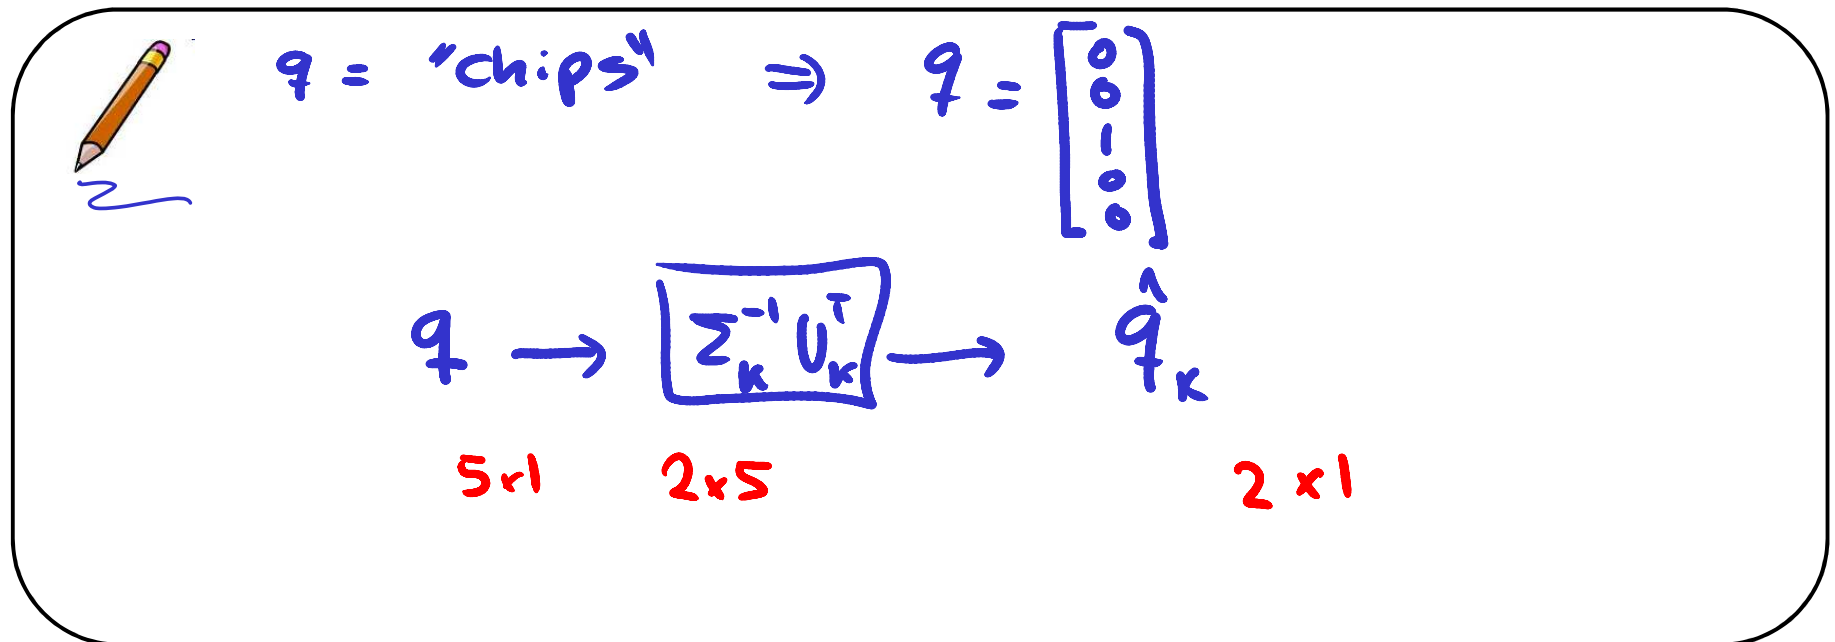
## Part II: Querying the search engine

To carry out **retrieval**, a **query**  $\mathbf{q} \in \mathbb{R}^n$  is first projected to the low-dimensional space:

$$\hat{\mathbf{q}}_k = \Sigma_k^{-1} \mathbf{U}_k^T \mathbf{q}$$

$2 \times 1$        $2 \times 5$        $5 \times 1$

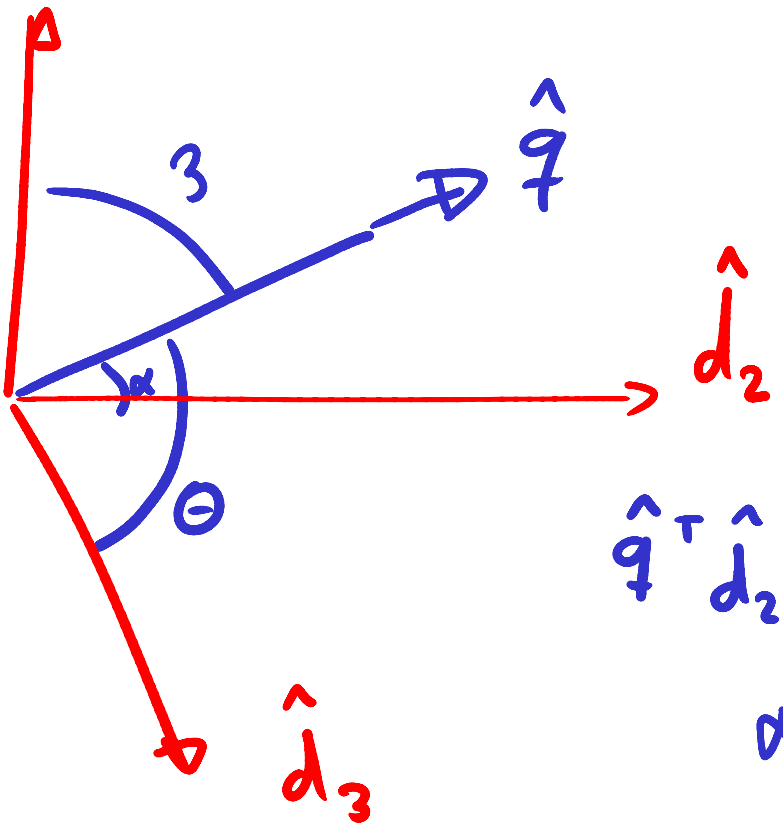
And then we measure the angle between  $\hat{\mathbf{q}}_k$  and the  $\mathbf{v}_k$ .



# Part II: Querying the search engine



$\hat{d}_1$



$$V^T = \begin{bmatrix} \hat{d}_1 & \hat{d}_2 & \hat{d}_3 \\ \circledast & \circledast & \circledast \\ \circledast & \circledast & \circledast \\ \circledast & \circledast & \circledast \end{bmatrix}$$

$$\hat{q}^T \hat{d}_2 = \|\hat{q}\| \|\hat{d}_2\| \cos \alpha$$

$$\alpha = \cos^{-1} \left( \frac{\hat{q}^T \hat{d}_2}{\|\hat{q}\| \|\hat{d}_2\|} \right)$$

# TF-IDF

The *term frequency* of word  $w$  in document  $d$  is equal to the total number of times  $w$  appears in  $d$ , divided by the total number of words in  $d$ .

The *inverse document frequency* of word  $w$  is based on the logarithm of the inverse frequency of the word in the corpus. If the number of documents in the corpus is  $D$  and the number of documents the word appears in is  $D_w$ , then

$$idf_w = \log \frac{D}{1 + D_w}$$

and for  $w$  and  $d$ , we can compute

$$tfidf_{w,d} = tf_{w,d}idf_w$$